

A generalisation of Nash's theorem with higher-order functionals

Julian Hedges

Queen Mary University of London, London E1 4NS, UK
jules.hedges@eecs.qmul.ac.uk

Abstract

The recent theory of sequential games and selection functions by Martin Escardó and Paulo Oliva is extended to games in which players move simultaneously. The Nash existence theorem for mixed-strategy equilibria of finite games is generalised to games defined by selection functions. A normal form construction is given which generalises the game-theoretic normal form, and its soundness is proven. Minimax strategies also generalise to the new class of games and are computed by the Berardi-Bezem-Coquand functional, studied in proof theory as an interpretation of the axiom of countable choice.

1 Introduction

The notion of *optimisation* is common to many areas of applied mathematics, such as game theory and linear and nonlinear programming. Typically we have a set X of *choices* and a function p mapping each $x \in X$ to a real number $p(x)$, which we might call the *value* or *cost* of x . From this we can define a natural notion of *optimality*: a point $y \in \mathbb{R}$ is optimal just if $y \geq p(x)$ for all $x \in X$ and $y = p(x_0)$ for some $x_0 \in X$. We usually refer to y by a notation such as

$$y = \max_{x \in X} p(x)$$

The point x_0 is also interesting: it is a point at which p *attains* its optimal value, and we refer to it as

$$x_0 = \arg \max_{x \in X} p(x)$$

(Of course, while y is guaranteed to be unique when it exists, x_0 is not necessarily unique; we only require that $\arg \max$ chooses some value for x_0 .) These notations are connected by the equation $y = p(x_0)$, or

$$\max_{x \in X} p(x) = p \left(\arg \max_{x \in X} p(x) \right)$$

Suppose we fix the set X and assume that $\max_{x \in X} p(x)$ exists for all functions $p : X \rightarrow \mathbb{R}$ (as when X is finite, for example). We can now define a function by

$$\varphi(p) = \max_{x \in X} p(x)$$

φ has range \mathbb{R} , and its domain is the function set $X \rightarrow \mathbb{R}$, that is, the set of all functions with domain X and range \mathbb{R} . We therefore write

$$\varphi : (X \rightarrow \mathbb{R}) \rightarrow \mathbb{R}$$

We call φ a *higher-order function*, that is, a function whose domain is itself a set of functions. We can also define

$$\varepsilon(p) = \arg \max_{x \in X} p(x)$$

obtaining a higher-order function

$$\varepsilon : (X \rightarrow \mathbb{R}) \rightarrow X$$

satisfying

$$\varphi(p) = p(\varepsilon(p)) \text{ for all } p : X \rightarrow \mathbb{R}$$

Using the concept of a higher-order function we can make a large generalisation of the properties of max and arg max. For any sets X and R , a function $\varphi : (X \rightarrow R) \rightarrow R$ will be called a *quantifier* and a function $\varepsilon : (X \rightarrow R) \rightarrow X$ will be called a *selection function*. We say that ε *attains* φ just if $\varphi(p) = p(\varepsilon(p))$ for all $p : X \rightarrow R$. max and arg max become the prototypical examples of a quantifier and a selection function attaining it. A very different example of a quantifier is a fixed point operator $\mu : (X \rightarrow X) \rightarrow X$ which has the property that $\mu(p)$ is always a fixed point of p , that is, $\mu(p) = p(\mu(p))$. Thus a fixed point operator attains itself. Quantifiers where R is the set of truth-values appear naturally in logic. These concepts were introduced and applied to the theory of sequential games by Martin Escardó and Paulo Oliva in a series of papers summarised in (Escardó and Oliva 2011).

What is a game? Typically, some players take turns choosing between sets of legal moves, which may be constrained by previous players' moves. The sequence of moves made by the players is called a *play* of the game. Usually, the rules of the game guarantee that every play terminates after a finite number of moves, and then uniquely determine which player has won the play.

In the theory of games as introduced by (von Neumann and Morgenstern 1944) the notion of a player winning a play is not used. Rather, for each player the rules of the game define an *outcome function* mapping each play of the game to a real number called the *utility* of the play for that player. This generalisation is important for applications of game theory to economics, where utility often represents profit. In the game played by two competing firms, for example, each firm is interested in maximising its own profit, and does not care (in the short term, at least) how much profit its competitor makes. Of course, a firm's profits will be affected by the moves of its competitor, and vice versa. A central problem of game theory is to determine which moves each player should choose in order to maximise their utility. The theory of games as surveyed for example in (Fudenberg and Tirole 1991) will be referred to as *classical game theory*.

Suppose during the course of a play some player must choose between some set X of moves. Taking the usual assumption of common knowledge of rationality (that is, the players play optimally, and they know that each other will play optimally, and so on) the future of the play after making each choice of $x \in X$ is sufficiently well determined that a utility $p(x) \in \mathbb{R}$ can be assigned

to each $x \in X$. In classical game theory, a rational player will always choose $\arg \max_{x \in X} p(x)$. By replacing \mathbb{R} with an arbitrary set R and $\arg \max$ with an arbitrary selection function $\varepsilon : (X \rightarrow R) \rightarrow X$, a rich theory of generalised games results, with deep connections to proof theory and theoretical computer science (Escardó and Oliva 2010b, Escardó et al. 2011).

The games which have been described so far are the so-called *sequential* games. In the more usual language of classical game theory this can be read as *non-branching extensive form games of perfect information*. However there are games which cannot be described as a sequence of moves. These are the so-called *simultaneous* games, or *games of imperfect information*. A well known example is rock-paper-scissors; a more important example is the simultaneous pricing of goods by supermarkets. von Neumann and Morgenstern proved that every game can be described as a simultaneous game, called its *normal* or *strategic form*. The central idea of this proof is that players simultaneously choose *contingent strategies*, higher-order functions which choose the next move given the play up to that point, and so play the game on behalf of the player. In this paper we consider a notion of simultaneous games that encompasses Escardó and Oliva’s generalised sequential games in a similar way.

In section 3, generalised simultaneous games and their appropriate notion of equilibrium are defined. In section 4 a class of games, the so-called *multilinear games*, is defined, and it is proven that games of this kind always have an equilibrium (theorem 2). This is used in section 5 to prove the key result of this paper (theorem 3), a natural generalisation of Nash’s theorem for the existence of mixed-strategy equilibria to games defined by arbitrary quantifiers. In section 6 a mapping from sequential to simultaneous games is defined analogous to the normal form construction in the classical theory, and its soundness is proven (theorem 5). In section 7 we show an interesting connection to proof theory, namely that the *binary Barardi-Bezem-Coquand functional* computes minimax strategies of games, a result that suggests a deeper connection between proof theory and generalised games.

2 Preliminaries

If X and Y are sets then $X \rightarrow Y$ denotes the set of all functions with domain X and range Y (this is often denoted Y^X , a notation we avoid in order to avoid writing exponential towers for higher-order functions). Cartesian products of sets are denoted \prod and bind tighter than \rightarrow , so for example $\prod_{i \in I} X_i \rightarrow R$ means $(\prod_{i \in I} X_i) \rightarrow R$. The i th coordinate projection of a tuple $\pi \in \prod_{i \in I} X_i$ is denoted π_i .

The following piece of notation, for manipulating products, will be helpful. Let I be a set and let X_i be a set for each $i \in I$. If $x \in X_i$ and $\pi \in \prod_{j \in I} X_j$ then we define $\pi(i \mapsto x) \in \prod_{j \in I} X_j$ by

$$(\pi(i \mapsto x))_j = \begin{cases} x & \text{if } i = j \\ \pi_j & \text{otherwise} \end{cases}$$

We make use of Church’s λ -notation for describing functions anonymously. The function which might otherwise be written as $x \mapsto 1 + x$ will be denoted $\lambda x^{\mathbb{N}}.1 + x$, where \mathbb{N} is the domain of the anonymous function. For example we

have $(\lambda x^{\mathbb{N}}.1 + x)(42) = 43$. A variable bound by a λ need not appear under the scope of the λ , for example $\lambda x^X.42$ is the constant function with the property that $(\lambda x^X.42)(x') = 42$ for all $x' \in X$.

A *quantifier* is a function $\varphi \in S_R(X)$ where $S_R(X) = (X \rightarrow R) \rightarrow \mathcal{P}(R)$, a definition introduced in (Escardó and Oliva 2011). The *domain* of a quantifier is

$$\text{dom}(\varphi) = \{p \in X \rightarrow R \mid \varphi(p) \neq \emptyset\}$$

A quantifier with $\text{dom}(\varphi) = X \rightarrow R$ will be called *total*.

A *selection function* is a function $\varepsilon \in J_R(X)$ where $J_R(X) = (X \rightarrow R) \rightarrow X$. Selection functions were first introduced in (Escardó and Oliva 2010a). The quantifier $\varphi \in S_R(X)$ is *attained* by the selection function $\varepsilon \in J_R(X)$ just if

$$p(\varepsilon(p)) \in \varphi(p)$$

for all $p \in \text{dom}(\varphi)$. This definition of attainment differs from Escardó and Oliva's, who require the condition to hold for all $p \in X \rightarrow R$. For a total quantifier (which are considered in section 5, and to which the main theorem applies) the two definitions coincide.

For example, if $R = \mathbb{R}$ and X is compact then the extreme value theorem (plus the axiom of choice) implies that the maximum quantifier

$$\varphi(p) = \begin{cases} \{\max_{x \in X} p(x)\} & \text{if } p \text{ is continuous} \\ \emptyset & \text{otherwise} \end{cases}$$

is attained. (Note that we need the axiom of choice to collect all the values into a single function.) A quantifier such as this whose values have cardinality at most 1 will be called *single-valued*. Since the definition of φ is clumsy we can use a new notation for single-valued quantifiers, such as

$$\varphi(p) = \max_{x \in X} p(x) \Big|_{p \text{ is continuous}}$$

We assume some point-set topology as covered, for example, in (Kelley 1955) and elementary properties of topological vector spaces (Conway 1990). All topological vector spaces are assumed to be T_1 throughout (this is no loss of generality because quotienting a topological vector space by the closure of $\{0\}$ always yields a Hausdorff space). For reference, a subset S of a real vector space is called *convex* iff for all $x, y \in S$ and $t \in [0, 1]$ we have $tx + (1 - t)y \in S$.

In part 4 we work with the class of *locally convex spaces*. The definition of a locally convex space is technical and not necessary for our purposes; beyond theorem 1 and lemma 2 we only need to know that every locally convex space is a topological vector space. Every normed vector space is locally convex; examples of locally convex spaces which are not normable include the spaces of smooth functions $C^\infty(\mathbb{R})$ and $C^\infty([0, 1])$ and the space of real-valued sequences \mathbb{R}^ω with convergence defined pointwise. Locally convex spaces are covered in detail in (Conway 1990).

A note on foundations. It is possible to define generalised sequential games over models other than classical set theory. Indeed, as explained in (Escardó and Oliva 2011) it is sometimes necessary to work in nonstandard

models, for example when considering unbounded sequential games (which are not considered in this paper). J_R is a (strong) monad and can be defined over any cartesian closed category (moreover the closely related $K_R(X) = (X \rightarrow R) \rightarrow R$, which contains the total single-valued quantifiers, is already well-known from programming language theory where it is called the *continuation monad*). The definitions of generalised simultaneous game and abstract Nash equilibrium could be formalised in a more general setting, but the proofs in section 4 use classical set theory in an essential way, so we find it easier to avoid foundational issues altogether and work entirely in classical set theory.

3 Generalised simultaneous games

In this section we define the objects studied in this paper, namely *generalised simultaneous games* and *generalised Nash equilibria*. The definition of a generalised simultaneous game comes from the classical definition of a normal-form game, but with the maximising behaviour of players replaced with a specified quantifier. For the general definition we do not require the number of players to be finite. The related notion of *generalised sequential game* will be defined in section 6.

Definition 1 (Generalised simultaneous game). *A generalised simultaneous game (with multiple outcome spaces), denoted simply game when not ambiguous, is a tuple*

$$\mathcal{G} = (I, (X_i, R_i, q_i, \varphi_i)_{i \in I})$$

where I is a nonempty set of players, and for each $i \in I$,

- X_i is a nonempty set of moves for player i ;
- R_i is a set of outcomes for player i ;
- $q_i \in S \rightarrow R_i$ is the outcome function for player i , where $S = \prod_{j \in I} X_j$ is the strategy space of \mathcal{G} ;
- $\varphi_i \in S_{R_i}(X_i)$ is the quantifier for player i .

We say that \mathcal{G} has a single outcome space if the R_i are equal and the q_i are equal. In this case \mathcal{G} is determined by a tuple

$$\mathcal{G} = (I, (X_i)_{i \in I}, R, q, (\varphi_i)_{i \in I})$$

where $q \in S \rightarrow R$.

An element $x \in X_i$ is called a strategy for player i for \mathcal{G} . A tuple $\pi \in S$ is called a strategy profile for \mathcal{G} . Throughout this paper the variables π , σ and τ will range over strategies of a game.

In general we need games with multiple outcome spaces to study simultaneous games, and in particular to recover the classical Nash theorem. However normal forms of generalised sequential games will always have single outcome space.

The appropriate notion of equilibrium of a generalised simultaneous game is called a *generalised Nash equilibrium*. Before making this definition, we first define some notation used throughout this paper. Firstly we define the family of

unilateral maps \mathcal{U}_q^i , which are used as a shorthand notation but, when considered as a higher-order functions, are also natural and interesting in their own right.

Definition 2 (Unilateral map). *Let I be a set, and for each $i \in I$ let X_i and R_i be sets. Let $q = (q_i)_{i \in I}$ be a family of maps such that each*

$$q_i \in \prod_{j \in I} X_j \rightarrow R_i$$

We define the i th unilateral map

$$\mathcal{U}_q^i \in \prod_{j \in I} X_j \rightarrow (X_i \rightarrow R_i)$$

by

$$\mathcal{U}_q^i(\pi)(x) = q_i(\pi(i \mapsto x))$$

Thus, the i th unilateral map computes the outcomes of unilateral changes of strategy by the i th player in a game. Secondly, we associate to every quantifier a set called its *diagonal*.

Definition 3 (Diagonal of a quantifier). *Let $\varphi \in S_R(X)$ be a quantifier. The diagonal of φ is*

$$\Delta(\varphi) = \{(p, x) \in (X \rightarrow R) \times X \mid p(x) \in \varphi(p)\}$$

Now the equilibria of a generalised simultaneous game can be defined in a very compact and (as will be seen) useful way.

Definition 4 (Generalised Nash equilibrium). *Let \mathcal{G} be a game with strategy space S . We define the best response correspondence $B \in S \rightarrow \mathcal{P}(S)$ of \mathcal{G} by*

$$B(\pi) = \bigcap_{i \in I} B_i(\pi)$$

where the $B_i \in S \rightarrow \mathcal{P}(S)$ are defined by

$$B_i(\pi) = \{\sigma \in S \mid (\mathcal{U}_q^i(\pi), \sigma_i) \in \Delta(\varphi_i)\}$$

A generalised Nash equilibrium of \mathcal{G} is a fixed point of B , that is, a strategy profile π such that $\pi \in B(\pi)$.

Unpacking this definition, we see that π is a generalised Nash equilibrium of \mathcal{G} iff for each $i \in I$ we have

$$q_i(\pi) \in \varphi_i(\lambda x^{X_i}. q_i(\pi(i \mapsto x)))$$

When X_i is compact, q_i is continuous and φ_i is the quantifier

$$\varphi_i(p) = \max_{x \in X_i} p(x) \Big|_p \text{ is continuous}$$

this reduces to

$$q_i(\pi) = \max_{x \in X_i} q_i(\pi(i \mapsto x))$$

which is the usual definition of a Nash equilibrium.

4 Multilinear games

Now we define a large family of games, called the *multilinear games*, that are guaranteed to have a generalised Nash equilibrium. The structure of the argument is the same as that in (Nash 1950b), but given in more generality to deal with more general quantifiers. This section can be seen as a series of lemmas that are eventually used to prove theorem 3 (the generalisation of Nash's theorem) in the next section.

Definition 5 (Closed graph property). *Let X and Y be topological spaces and $F \in X \rightarrow \mathcal{P}(Y)$. We say that F has the closed graph property iff*

$$\Gamma(F) = \{(x, y) \in X \times Y \mid y \in F(x)\}$$

is closed with respect to the product topology.

The closed graph property is a form of continuity for functions whose range is a set of subsets of a topological space.

In order to guarantee that a generalised simultaneous game will have an equilibrium we need to impose closed graph properties on the quantifiers. However the domain of a quantifier is a function set, which in general has no unique natural topology. The least we need is that the unilateral maps are continuous, and so for this reason we define the *unilateral topology*.

Definition 6 (Unilateral topology). *For each $i \in I$ let X_i and R_i be topological spaces with $q_i \in X_i \rightarrow R_i$ continuous. The unilateral topology on $X_i \rightarrow R_i$ is the final topology with respect to the singleton family $\{\mathcal{U}_q^i\}$, that is, it is the largest topology with respect to which \mathcal{U}_q^i is continuous. A function which is continuous with respect to the unilateral topology will be called *unilaterally continuous*, and a function which has closed graph with respect to the unilateral topology has *unilaterally closed graph*.*

Another possible topology on $X_i \rightarrow R_i$ which will be useful is the topology of pointwise convergence. Most of this paper could be formulated using only pointwise convergence, except for an interesting example at the end of this section which needs a finer topology, namely uniform convergence.

Lemma 1. *The unilateral topology is finer than the topology of pointwise convergence.*

Proof. It must be proven that \mathcal{U}_q^i is continuous with respect to the topology of pointwise convergence. Let $\pi_j \rightarrow \pi$ be a convergent sequence in $\prod_{j \in I} X_j$, and let $x \in X_i$. We have

$$\pi_j(i \mapsto x) \rightarrow \pi(i \mapsto x)$$

in the product topology, so

$$\mathcal{U}_q^i(\pi_j)(x) = q_i(\pi_j(i \mapsto x)) \rightarrow q_i(\pi(i \mapsto x)) = \mathcal{U}_q^i(\pi)(x)$$

because q_i is continuous. Therefore $\mathcal{U}_q^i(\pi_j) \rightarrow \mathcal{U}_q^i(\pi)$ pointwise, as required. \square

Now we can give the definition of a multilinear game. This definition essentially contains the least assumptions needed for Nash's proof.

Definition 7 (Multilinear game). A game $\mathcal{G} = (I, (X_i, R_i, q_i, \varphi_i)_{i \in I})$ is called multilinear iff

- Each X_i is a compact and convex subset of a given locally compact space V_i over \mathbb{R} ;
- Each R_i is a topological vector space over \mathbb{R} ;
- Each q_i extends to a continuous multilinear map

$$q_i \in \prod_{j \in I} V_j \rightarrow R_i$$

(that is, q_i is linear with respect to each V_j separately);

- Each φ_i has unilaterally closed graph, $\varphi_i(p)$ is closed and convex for all $p \in X_i \rightarrow R_i$, and $\text{dom}(\varphi_i) \supseteq \text{im}(\mathcal{U}_q^i)$.

(Note that because q_i is continuous and multilinear, to satisfy the last condition it suffices that $\varphi_i(p) \neq \emptyset$ whenever p is continuous and linear. Note also that if φ_i is single-valued then each $\varphi_i(p)$ is automatically closed and convex.)

The idea of the existence proof is to reduce to the following fixed point theorem.

Theorem 1 (Kakutani-Fan-Glicksberg fixed point theorem (Fan 1952, Glicksberg 1952)). Let S be a nonempty, compact and convex subset of a locally convex space over \mathbb{R} . Let $B \in S \rightarrow \mathcal{P}(S)$ have closed graph and let $B(\pi)$ be nonempty, closed and convex for all $\pi \in S$. Then B has a fixed point.

We will need to use the fact that locally convex spaces are closed under arbitrary products.

Lemma 2. Let $\{V_i\}_{i \in I}$ be a family of locally convex spaces over a field K . Then $\prod_{i \in I} V_i$ has the structure of a locally convex space over K whose topology is the product topology.

Much of the usefulness of multilinear games comes down to the fact that their unilateral maps are well-behaved.

Lemma 3. Let \mathcal{G} be a multilinear game with strategy space S . Then each \mathcal{U}_q^i is a continuous function

$$\mathcal{U}_q^i \in S \times X_i \rightarrow R_i$$

(under the Curry bijection $A \rightarrow (B \rightarrow C) \cong A \times B \rightarrow C$) which is linear in its second argument.

Proof. By the continuity and multilinearity of the q_i . □

Lemmas 4-8 form the core of the proof, establishing the hypotheses of the Kakutani-Glicksberg-Fan theorem.

Lemma 4. Let \mathcal{G} be a multilinear game. Then the strategy space of \mathcal{G} is a nonempty, compact and convex subset of a locally convex space.

Proof. The strategy space is

$$S = \prod_{i \in I} X_i \subseteq \prod_{i \in I} V_i$$

where the larger space is locally convex by lemma 2.

S is nonempty by the axiom of choice and compact by Tychonoff's theorem. Convexity is also inherited by the product, since for each $i \in I$ we have

$$(tx + (1-t)y)_i = tx_i + (1-t)y_i \in X_i \quad \square$$

Lemma 5. *Let \mathcal{G} be a multilinear game with best response correspondence B such that each quantifier φ_i is attained by a selection function ε_i . Then $B(\pi)$ is nonempty for all π .*

Proof. Let S be the strategy space of \mathcal{G} . Given $\pi \in S$ we define $\sigma \in S$ to have i th component

$$\sigma_i = \varepsilon_i(\mathcal{U}_q^i(\pi))$$

Since ε_i attains φ_i and $\mathcal{U}_q^i(\pi) \in \text{dom}(\varphi_i)$ we have

$$\mathcal{U}_q^i(\pi)(\varepsilon_i(\mathcal{U}_q^i(\pi))) \in \varphi_i(\mathcal{U}_q^i(\pi))$$

Therefore

$$(\mathcal{U}_q^i(\pi), \sigma_i) = (\mathcal{U}_q^i(\pi), \varepsilon_i(\mathcal{U}_q^i(\pi))) \in \Delta(\varphi_i)$$

so $\sigma \in B(\pi)$, as required. \square

Lemma 6. *Let \mathcal{G} be a multilinear game with best response correspondence B . Then $B(\pi)$ is closed for all π .*

Proof. It suffices to prove that each factor

$$B_i(\pi) = \{\sigma \in S \mid (\mathcal{U}_q^i(\pi), \sigma_i) \in \Delta(\varphi_i)\}$$

is closed. Let $\sigma_j \rightarrow \sigma$ be a convergent sequence in $B_i(\pi)$. For each i we have $\sigma_{j,i} \rightarrow \sigma_i$, so

$$\mathcal{U}_q^i(\pi)(\sigma_{j,i}) \rightarrow \mathcal{U}_q^i(\pi)(\sigma_i)$$

by the continuity of \mathcal{U}_q^i . We also have that each

$$\mathcal{U}_q^i(\pi)(\sigma_{j,i}) \in \varphi_i(\mathcal{U}_q^i(\pi))$$

and the right hand side is closed by definition, therefore

$$\mathcal{U}_q^i(\pi)(\sigma_i) \in \varphi_i(\mathcal{U}_q^i(\pi))$$

that is,

$$(\mathcal{U}_q^i(\pi), \sigma_i) \in \Delta(\varphi_i) \quad \square$$

Lemma 7. *Let \mathcal{G} be a multilinear game with best response correspondence B . Then $B(\pi)$ is convex for all π .*

Proof. Suppose $\sigma, \tau \in B(\pi)$ and $t \in [0, 1]$. Let $i \in I$. By definition we have

$$\mathcal{U}_q^i(\pi)(\sigma_i), \mathcal{U}_q^i(\pi)(\tau_i) \in \varphi_i(\mathcal{U}_q^i(\pi))$$

Since the linearity of \mathcal{U}_q^i we have

$$\mathcal{U}_q^i(\pi)(t\sigma_i + (1-t)\tau_i) = t\mathcal{U}_q^i(\pi)(\sigma_i) + (1-t)\mathcal{U}_q^i(\pi)(\tau_i)$$

Since the $\varphi_i(p)$ are convex, we have

$$\mathcal{U}_q^i(\pi)(t\sigma_i + (1-t)\tau_i) \in \varphi_i(\mathcal{U}_q^i(\pi))$$

that is,

$$(\mathcal{U}_q^i(\pi), t\sigma_i + (1-t)\tau_i) \in \Delta(\varphi_i)$$

Therefore

$$t\sigma + (1-t)\tau \in B(\pi) \quad \square$$

Lemma 8. *Let \mathcal{G} be a multilinear game with best response correspondence B . Then B has closed graph.*

Proof. Note that

$$\Gamma(B) = \bigcap_{i \in I} \{(\sigma, \pi) \in S^2 \mid (\mathcal{U}_q^i(\pi), \mathcal{U}_q^i(\pi)(\sigma_i)) \in \Gamma(\varphi_i)\}$$

and so it suffices to prove these factors closed. Let $(\sigma_j, \pi_j) \rightarrow (\sigma, \pi)$ be a convergent sequence in the i th factor. By the continuity of \mathcal{U}_q^i ,

$$\mathcal{U}_q^i(\pi_j)(\sigma_{j,i}) \rightarrow \mathcal{U}_q^i(\pi)(\sigma_i)$$

Since \mathcal{U}_q^i is also unilaterally continuous as a map $S \rightarrow (X_i \rightarrow R_i)$, we have $\mathcal{U}_q^i(\pi_j) \rightarrow \mathcal{U}_q^i(\pi)$ unilaterally. Therefore we have a convergent sequence

$$(\mathcal{U}_q^i(\pi_j), \mathcal{U}_q^i(\pi_j)(\sigma_{j,i})) \rightarrow (\mathcal{U}_q^i(\pi), \mathcal{U}_q^i(\pi)(\sigma_i))$$

in the graph $\Gamma(\varphi_i)$, which is closed by definition. \square

Theorem 2 (Existence theorem for multilinear games). *Let \mathcal{G} be a multilinear game such that each quantifier is attained by a selection function. Then \mathcal{G} has a generalised Nash equilibrium.*

Proof. Let B be the best response correspondence of \mathcal{G} . By lemmas 4-8 and the Kakutani-Fan-Glicksberg fixed point theorem, B a fixed point. \square

Examples of multilinear games as mixed extensions of finite games are given in the next section. Another interesting example is given by integration. Let $X_i = [0, 1]$, $V_i = \mathbb{R}$ and $R_i = \mathbb{R}$, and let $L(X_i)$ be the set of all Lebesgue-integrable functions $p \in [0, 1] \rightarrow \mathbb{R}$ with

$$\left| \int_0^1 p(x) dx \right| < \infty$$

Define a single-valued quantifier $\varphi_i \in S_{\mathbb{R}} X_i$ by

$$\varphi_i(p) = \int_0^1 p(x) dx \Big|_{p \in L(X_i)}$$

Using the mean value theorem (and the axiom of choice) we can prove the existence of a selection function attaining φ_i : for all $p \in L(X_i)$ there exists $\varepsilon_i(p) \in X_i$ such that

$$p(\varepsilon_i(p)) = \int_0^1 p(x) dx$$

This highly nonconstructive selection function was briefly introduced as an example in (Escardó and Oliva 2010a).

We let I be finite and for simplicity let the other X_j be normed, so the strategy space is normed and we can work with the $\varepsilon - \delta$ definitions of uniform convergence and continuity.

Lemma 9. *If q_i is uniformly continuous and $\pi_j \rightarrow \pi$ then $\mathcal{U}_q^i(\pi_j) \rightarrow \mathcal{U}_q^i(\pi)$ uniformly.*

Proof. We have that q_i is uniformly continuous, that is,

$$\forall \varepsilon > 0 \exists \delta > 0 \forall \pi, \sigma \in \prod_{j \in I} X_j. |\pi - \sigma| < \delta \implies |q_i(\pi) - q_i(\sigma)| < \varepsilon \quad (1)$$

We also have $\pi_j \rightarrow \pi$, that is,

$$\forall \varepsilon > 0 \exists N \forall j \geq N. |\pi_j - \pi| < \varepsilon \quad (2)$$

We want to prove that $\mathcal{U}_q^i(\pi_j) \rightarrow \mathcal{U}_q^i(\pi)$ uniformly, that is,

$$\forall \varepsilon > 0 \exists N \forall j \geq N \forall x \in X_i. |\mathcal{U}_q^i(\pi_j)(x) - \mathcal{U}_q^i(\pi)(x)| < \varepsilon$$

Let $\varepsilon > 0$. By (1), we have $\delta > 0$ with the given property. We take ε in (2) to be this δ , obtaining N . Let $j \geq N$, therefore

$$|\pi_j - \pi| < \delta$$

by (2). Let $x \in X$. The crucial observation is that $\pi_j(i \mapsto x)$ behaves like π_j but is constant in its i th coordinate. That is, we have

$$|\pi_j(i \mapsto x) - \pi(i \mapsto x)| \leq |\pi_j - \pi| < \delta$$

Now we take π, σ in (1) to be $\pi_j(i \mapsto x)$ and $\pi(i \mapsto x)$. We have already proved the antecedent in (1), therefore

$$|\mathcal{U}_q^i(\pi_j)(x) - \mathcal{U}_q^i(\pi)(x)| = |q_i(\pi_j(i \mapsto x)) - q_i(\pi(i \mapsto x))| < \varepsilon$$

as required. \square

We have proven that the unilateral topology is finer than the topology of uniform convergence.

Lemma 10. *φ_i is unilaterally continuous.*

Proof. Suppose we have $p_j \rightarrow p$ uniformly in $L(X_i)$. Since the convergence of the integrands is uniform, we can apply the uniform convergence theorem to get

$$\varphi_i(p_j) = \int_0^1 p_j(x) dx \rightarrow \int_0^1 p(x) dx = \varphi_i(p)$$

Since the unilateral topology is finer than the topology of uniform convergence, we are done. \square

In the 1-player game defined by the integration quantifier with outcome function q , the unique value of $q(x)$ when x is an equilibrium strategy, which can be called the *expected outcome* of the game, is simply

$$\int_0^1 q(x) dx$$

In the 2-player game with both quantifiers integrals, a generalised Nash equilibrium (a, b) satisfies

$$a = \int_0^1 q_X(x, b) dx \quad b = \int_0^1 q_Y(a, y) dy$$

Since $\int_0^1 p(x) dx$ is the *average* value of p , this is a game where players are trying to gain the average outcome rather than the maximum. The existence of a Nash equilibrium in this case can be more directly proven by applying the Brouwer fixed point theorem to the mapping

$$[0, 1]^2 \rightarrow [0, 1]^2, (a, b) \mapsto \left(\int_0^1 q_X(x, b) dx, \int_0^1 q_Y(a, y) dy \right)$$

5 Finite games

In this section we apply the existence theorem for multilinear games to prove a suitable generalisation of Nash's theorem. The classical version of Nash's theorem guarantees that every finite game (that is, a classical game in which each player has finitely many strategies) has a *mixed strategy Nash equilibrium*.

The notion of mixed strategies means that we consider probability distributions over ordinary strategies (referred to as *pure strategies* for clarity). The outcome functions also need to be replaced by *expected outcome* functions. However the discussion of probability distributions can be avoided by treating them as geometric objects, namely *simplices*. This approach also makes it clearer how quantifiers and selection functions must be modified when passing to mixed strategies. A probabilistic interpretation of the resulting theorem is possible, but is avoided in this paper.

Definition 8 (Finite game). *A game $\mathcal{G} = (I, (X_i, R_i, q_i, \varphi_i)_{i \in I})$ is called finite iff*

- *I is finite;*
- *Each X_i is finite;*
- *Each R_i is a topological vector space over \mathbb{R} ;*

- Each φ_i is total, has closed graph with respect to the topology of pointwise convergence, and $\varphi_i(p)$ is closed and convex for all $p \in X_i \rightarrow R_i$.

Note that restricting to pointwise convergence is no loss of generality here because the X_i are finite.

The set of probability distributions over a finite set can be seen as a geometric object called a *standard simplex*. In 2 and 3 dimensions these can be easily visualised as a line segment and an equilateral triangle; the next simplex Δ_4 is a tetrahedron seen as a subset of \mathbb{R}^4 .

Definition 9 (Standard simplex). *The n th standard simplex is the set*

$$\Delta_n = \{(x_1, \dots, x_n) \in \mathbb{R}^n \mid \sum_{i=1}^n x_i = 1 \text{ and each } x_i \geq 0\}$$

Definition 10 (Mixed extension). *Let $\mathcal{G} = (I, (X_i, R_i, q_i, \varphi_i)_{i \in I})$ be a finite game with strategy space S . We define a game*

$$\mathcal{G}^* = (I, (X_i^*, R_i, q_i^*, \varphi_i^*)_{i \in I})$$

called the mixed extension of \mathcal{G} as follows: player i has move set

$$X_i^* = \Delta_{|X_i|}$$

outcome function

$$q_i^*(\pi) = \sum_{\sigma \in S} \left(\prod_{i \in I} \pi_{i, \sigma_i} \right) (q_i(\sigma))$$

and quantifier

$$\varphi_i^*(p) = \varphi_i(p \circ \delta_i)$$

where δ_i is the canonical injection $X_i \hookrightarrow X_i^$ mapping each j to the vertex of the simplex at which the j th coordinate is 1.*

Note that the finiteness of I is used only in the well-definition of the q_i^* : for the strategy space S to be finite it is necessary that I be finite, except in trivial cases when all but finitely many X_i are singletons.

Definition 11 (Mixed strategy abstract Nash equilibrium). *Let \mathcal{G} be a finite game. A strategy profile for \mathcal{G}^* will be called a mixed strategy profile for \mathcal{G} . An abstract Nash equilibrium of \mathcal{G}^* will be called a mixed strategy abstract Nash equilibrium of \mathcal{G} .*

The most important property of mixed extensions is that they are always multilinear. This will be used to prove the generalised Nash theorem.

Lemma 11. *Let \mathcal{G} be a finite game. Then \mathcal{G}^* is a multilinear game.*

Proof. Each Δ_n for $n > 0$ is a nonempty, compact and convex subset of the

locally convex space \mathbb{R}^n . Continuity of the q_i^* is clear. q_i^* is multilinear because

$$\begin{aligned}
& \mathcal{U}_{q^*}^i(\pi)(cx + y) \\
&= q^*(\pi(i \mapsto cx + y)) \\
&= \sum_{\sigma \in S} \left(\left(\prod_{j \neq i} \pi_{j, \sigma_j} \right) \cdot (cx_{\sigma_i} + y_{\sigma_i}) \right) (q(\sigma)) \\
&= c \sum_{\sigma \in S} \left(\left(\prod_{j \neq i} \pi_{j, \sigma_j} \right) \cdot x_{\sigma_i} \right) (q(\sigma)) + \sum_{\sigma \in S} \left(\left(\prod_{j \neq i} \pi_{j, \sigma_j} \right) \cdot y_{\sigma_i} \right) (q(\sigma)) \\
&= cq^*(\pi(i \mapsto x)) + q^*(\pi(i \mapsto y)) \\
&= c\mathcal{U}_{q^*}^i(\pi)(x) + \mathcal{U}_{q^*}^i(\pi)(y)
\end{aligned}$$

The $\varphi_i^*(p)$ are of the form $\varphi_i(p')$, and so are closed and convex. We note that φ_i^* is total because φ_i is. The graph of φ_i^* is

$$\Gamma(\varphi_i^*) = \{(p, y) \in (X_i^* \rightarrow R_i) \times R_i \mid (p \circ \delta_i, y) \in \Gamma(\varphi_i)\}$$

Suppose we have a convergent sequence $(p_j, y_j) \rightarrow (p, y)$ in $\Gamma(\varphi_i^*)$. Let $x \in X$, then

$$(p_j \circ \delta_i)(x) = p_j(\delta_i(x)) \rightarrow p(\delta_i(x)) = (p \circ \delta_i)(x)$$

Therefore $p_j \circ \delta_i \rightarrow p \circ \delta_i$ pointwise, so

$$(p_j \circ \delta_i, y_j) \rightarrow (p \circ \delta_i, y)$$

with respect to the topology of pointwise convergence. Since the unilateral topology is finer than the topology of pointwise convergence, we are done. \square

The final result we need is the ability to lift selection functions to mixed extensions.

Lemma 12. *Let X be a nonempty finite set and let $\varphi \in S_{\mathbb{R}}X$ be a total quantifier attained by the selection function $\varepsilon \in J_{\mathbb{R}}X$. Then there exists a selection function ε^* such that φ^* is attained by ε^* .*

Proof. We define $\varepsilon^* \in J_{\mathbb{R}}\Delta_{|X|}$ by the equation

$$\varepsilon^*(p) = \delta(\varepsilon(p \circ \delta))$$

where $\delta : X \hookrightarrow \Delta_{|X|}$. Then

$$p(\varepsilon^*(p)) = p(\varepsilon(p \circ \delta)) = (p \circ \delta)(\varepsilon(p \circ \delta)) \in \varphi(p \circ \delta) = \varphi^*(p)$$

for all $p \in \Delta_{|X|} \rightarrow \mathbb{R}$. \square

Theorem 3 (Existence theorem for finite games). *Let \mathcal{G} be a finite game such that each φ_i is attained by a selection function. Then \mathcal{G} has a mixed strategy abstract Nash equilibrium.*

Proof. \mathcal{G}^* is a multilinear game which is attained by selection functions by lemmas 11 and 12. Therefore \mathcal{G}^* has an abstract Nash equilibrium by theorem 2. \square

In order to recover the classical Nash theorem we simply consider finite games whose outcome spaces are \mathbb{R} and define $q_i(\pi) \in \mathbb{R}$ to be the utility of π for player i , taking all quantifiers to be max. We could instead define a finite game with single outcome space \mathbb{R}^n and let $(q(\pi))_i$ be the utility of π for player i , and consider selection functions $\varepsilon_i \in J_{\mathbb{R}^n} X_i$ maximising the i th coordinate:

$$\varepsilon_i(p) = \arg \max_{x \in X_i} (p(x))_i$$

However the quantifiers attained by these quantifiers are continuous only if $n = 1$. This game has the same equilibria as the equivalent game with multiple outcome spaces, but the Nash theorem cannot be proven in this way. It is for this reason that we need to consider games with multiple outcome spaces, in contrast to generalised sequential games (which do not require continuity).

For a different example of a quantifier in a finite game, let R_i be normed and fix $\varepsilon > 0$ and $x_0 \in X_i$. Define

$$\varphi_i(p) = B_\varepsilon(p(x_0))$$

that is, the closed ε -ball around $p(x_0)$. This quantifier is attained by the constant selection function $\varepsilon(p) = x_0$. For a sequential game this would force the game to be trivial, but this is not the case here: for example, if φ_X is the quantifier defined here and φ_Y is the maximum quantifier with $R_Y = \mathbb{R}$ then a Nash equilibrium is a point (a, b) such that

$$|q_X(a, b) - q_X(x_0, b)| < \varepsilon \quad q_Y(a, b) = \max_{y \in Y} q_Y(a, y)$$

6 The normal form of a sequential game

In classical game theory every game can be put into the form of a simultaneous game called its *normal form*. The major motivation for defining generalised simultaneous games was to generalise this operation to give a notion of normal form for generalised sequential games. This construction is given in this section and a form of soundness of proven, namely that the solution concept for a generalised sequential game, the so-called *optimal strategies*, are mapped to generalised Nash equilibria.

Definition 12 (Generalised sequential game). *A generalised sequential game is determined by a set R of outcomes, a set X_i of moves and a quantifier $\varphi_i \in S_R X_i$ for each $1 \leq i \leq n$, and an outcome function $q \in \prod_{i=1}^n X_i \rightarrow R$. A strategy in a sequential game is a tuple*

$$\pi \in \prod_{i=1}^n \left(\prod_{j=1}^{i-1} X_j \rightarrow X_i \right)$$

The strategy π is called optimal iff for all $\vec{a} = (a_1, \dots, a_{i-1}) \in \prod_{j=1}^{i-1} X_j$ (where $i > 0$) we have

$$q(\vec{a}, b_i^{\vec{a}}, \dots, b_n^{\vec{a}}) \in \varphi_i(\lambda x^{X_i}. q(\vec{a}, x, b_{k+i}^{\vec{a}, x}, \dots, b_n^{\vec{a}, x}))$$

where

$$b_j^{\vec{a}} = \pi_j(\vec{a}, b_i^{\vec{a}}, \dots, b_{j-1}^{\vec{a}})$$

Given a strategy π in a game, its strategic play is $\pi^\dagger \in \prod_{i=1}^n X_i$ given by

$$\begin{aligned}\pi_1^\dagger &= \pi_1 \text{ (modulo the isomorphism } \prod_{j=1}^0 X_j \rightarrow X_1 = \{0\} \rightarrow X_1 \cong X_1) \\ \pi_{i+1}^\dagger &= \pi_{i+1}(\pi_1^\dagger, \dots, \pi_i^\dagger)\end{aligned}$$

The strategic play of an optimal strategy is called an optimal play.

To be precise, this notion of sequential game is called a *finite game with multiple optimal outcomes* in (Escardó and Oliva 2011). Infinite games are avoided in this paper for simplicity.

Generalised sequential games were introduced in order to study a particular higher-order function called the *product of selection functions*. This is the function

$$\otimes \in J_R X \times J_R Y \rightarrow J_R(X \times Y)$$

given by

$$(\varepsilon \otimes \delta)(q) = (a, b_a)$$

where

$$a = \varepsilon(\lambda x^X. q(x, b_x)) \quad b_x = \delta(\lambda y^Y. q(x, y))$$

The product of selection functions has many interesting and unintuitive properties, especially when infinitely iterated: for example, it computes witnesses for the axiom of countable choice, and computes exhaustive searches of certain infinite types in finite time (Escardó and Oliva 2010b), both of which popular belief would have is impossible. Every use of the product of selection functions can be seen as the computation of an optimal play for a suitable generalised sequential game.

Theorem 4. *Let \mathcal{G} be a generalised sequential game whose quantifiers are total and attained by selection functions $\varepsilon_i \in J_R X_i$. Then*

$$\left(\bigotimes_{i=1}^n \varepsilon_i \right) (q)$$

is an optimal play for \mathcal{G} .

Now we give the normal form construction and prove that it maps optimal strategies to generalised Nash equilibria.

Definition 13 (Normal form). *Let \mathcal{G} be a generalised sequential game. We define a simultaneous game with single outcome space*

$$\mathcal{G}^\dagger = (I^\dagger, (X_i^\dagger)_{i \in I}, R, q^\dagger, (\varphi_i^\dagger)_{i \in I})$$

called the normal form of \mathcal{G} as follows:

- $I^\dagger = \{1, \dots, n\}$;
- Each $X_i^\dagger = P_i \rightarrow X_i$ where $P_i = \prod_{j=1}^{i-1} X_j$;
- $q^\dagger(\pi) = q(\pi^\dagger)$;

- Each $\varphi_i^\dagger(p) = \varphi_i(\lambda x^{X_i}.p(\lambda \sigma^{P_i}.x))$.

Theorem 5. Let \mathcal{G} be a sequential game and let π be an optimal strategy for \mathcal{G} . Then π is a generalised Nash equilibrium of \mathcal{G}^\dagger .

Proof. Let $1 \leq i \leq n$. It must be proven that

$$(\mathcal{U}_{q^\dagger}^i(\pi), \pi_i) \in \Delta(\varphi_i^\dagger)$$

Let $\vec{a} = (\pi_{j=1}^{i-1})^\dagger \in \prod_{j=1}^{i-1} X_j$. Since π is an optimal strategy for \mathcal{G} we have

$$q(\vec{a}, b_i^{\vec{a}}, \dots, b_n^{\vec{a}}) \in \varphi_i(\lambda x^{X_i}.q(\vec{a}, x, b_{i+1}^{\vec{a},x}, \dots, b_n^{\vec{a},x}))$$

By induction on j we have

$$b_j^{\vec{a}} = \pi_j(\vec{a}, b_i^{\vec{a}}, \dots, b_{j-1}^{\vec{a}}) = \pi_j(\pi_1^\dagger, \dots, \pi_{i-1}^\dagger, \pi_i^\dagger, \dots, \pi_{j-1}^\dagger) = \pi_j^\dagger$$

therefore

$$q^\dagger(\pi) = q(\pi^\dagger) = q(\vec{a}, b_i^{\vec{a}}, \dots, b_n^{\vec{a}})$$

We also have

$$\begin{aligned} & \varphi_i^\dagger(\mathcal{U}_{q^\dagger}^i(\pi)) \\ &= \varphi_i(\lambda x^{X_i}.\mathcal{U}_{q^\dagger}^i(\pi)(\lambda \sigma^{P_i}.x)) && \text{(definition of } \varphi_i^\dagger) \\ &= \varphi_i(\lambda x^{X_i}.q^\dagger(\pi(i \mapsto \lambda \sigma^{P_i}.x))) && \text{(definition of } \mathcal{U}_{q^\dagger}^i) \\ &= \varphi_i(\lambda x^{X_i}.q(\tau^\dagger)) && \text{(definition of } q^\dagger) \end{aligned}$$

where $\tau = \pi(i \mapsto \lambda \sigma^{P_i}.x)$. By induction on j we have

$$\tau_j^\dagger = \begin{cases} \pi_j^\dagger & \text{if } 1 \leq j < i \\ x & \text{if } i = j \\ \pi_j(\pi_1^\dagger, \dots, \pi_{i-1}^\dagger, x, \tau_{i+1}^\dagger, \dots, \tau_{j-1}^\dagger) & \text{if } i < j \leq n \end{cases}$$

We certainly have that τ^\dagger coincides with $(\vec{a}, x, b_{i+1}^{\vec{a},x}, \dots, b_n^{\vec{a},x})$ at indices $1 \leq j \leq i$. Moreover by induction on $i < j \leq n$ we have

$$b_j^{\vec{a},x} = \pi_j(\vec{a}, x, b_{i+1}^{\vec{a},x}, \dots, b_{j-1}^{\vec{a},x}) = \pi_j(\vec{a}, x, \tau_{i+1}^\dagger, \dots, \tau_{j-1}^\dagger) = \tau_j^\dagger$$

therefore

$$\tau^\dagger = (\vec{a}, x, b_{i+1}^{\vec{a},x}, \dots, b_n^{\vec{a},x})$$

We have therefore proven

$$q^\dagger(\pi) \in \varphi_i^\dagger(\mathcal{U}_{q^\dagger}^i(\pi))$$

that is,

$$(\mathcal{U}_{q^\dagger}^i(\pi), \pi_i) \in \Delta(\varphi_i^\dagger) \quad \square$$

The converse is false because optimal strategies of sequential games generalise the classical notion of subgame-perfect equilibrium, which is a stronger condition than classical Nash equilibrium (called an *equilibrium refinement* in classical game theory) (Escardó and Oliva 2012).

7 2-player games and minimax strategies

In this section the abstract notion of a ψ - φ strategy is defined, and used to show an intriguing connection between generalised simultaneous games and proof theory. The reason for this terminology is that a ψ - φ strategy corresponds to a minimax strategy in a 2-player game with $\varphi = \max$ and $\psi = \min$. Note however that when modelling a classical game as a generalised game all the quantifiers will be max, and so ψ - φ strategies are distinct in this sense from minimax strategies.

Definition 14 (ψ - φ strategy). *Let \mathcal{G} be a 2-player game with quantifiers $\varphi \in S_{R_X}X$ and $\psi \in S_{R_Y}Y$. A strategy $a \in X$ is called a ψ - φ strategy for player 1 iff*

$$q_1(a, f(a)) \in \varphi(\lambda x^X. q_1(x, f(x)))$$

for all $f \in X \rightarrow Y$ with the property that for all $x \in X$,

$$q_2(x, f(x)) \in \psi(\lambda y^Y. q_2(x, y))$$

Similarly b is a ψ - φ strategy for player 2 iff $q_2(g(b), b) \in \psi(\lambda y^Y. q_2(g(y), y))$ whenever $q_1(g(y), y) \in \varphi(\lambda x^X. q_1(x, y))$. A ψ - φ strategy profile is one whose components are both ψ - φ .

The binary Berardi-Bezem-Coquand functional is the higher-order function

$$\widehat{\otimes} \in J_R X \times J_R Y \rightarrow J_R(X \times Y)$$

given by

$$(\varepsilon \widehat{\otimes} \delta)(q) = (a, b)$$

where

$$\begin{aligned} a &= \varepsilon(\lambda x^X. q(x, \delta(\lambda y^Y. q(x, y)))) \\ b &= \delta(\lambda y^Y. q(\varepsilon(\lambda x^X. q(x, y)), y)) \end{aligned}$$

Notice that the type of $\widehat{\otimes}$ is the same as the type of \otimes . Moreover when infinitely iterated both provide proof interpretations (in the modified realizability interpretation of Heyting arithmetic) of the axiom of countable choice (Berardi et al. 1998, Berger 2002) and a certain equivalence, namely interdefinability over system T, is shown in (Powell 2012). However the relationship between the two functionals is not well understood, and only the product of selection functions has previously been linked to game theory.

Theorem 6. *Let \mathcal{G} be a 2-player game with single outcome space such and total single-valued quantifiers attained by selection functions ε, δ . Then the strategy profile*

$$(\varepsilon \widehat{\otimes} \delta)(q) \in X \times Y$$

is a ψ - φ strategy profile.

Proof. Since ψ is single-valued, the unique f with the given property is

$$f(x) = b_x = \delta(\lambda y^Y. q(x, y))$$

The first component in the strategy profile is

$$a = \varepsilon(\lambda x^X . q(x, b_x))$$

We therefore have

$$q(a, b_a) = (\lambda x^X . q(x, b_x))(\varepsilon(\lambda x^X . q(x, b_x))) \in \varphi(\lambda x^X . q(x, b_x))$$

as required. The proof for b is symmetric. \square

In particular, in a 2-player classical game if a player has outcome function q then they have a minimax strategy given by $(\arg \max \hat{\otimes} \arg \min)(q)$.

8 Conclusions

For sequential games, the proof of the existence of equilibria uses the product of selection functions and is constructive (theorem 4). Due to the importance of the product of selection functions, the constructive nature of the existence proof is an important part of the theory. Theorem 3 is similar to theorem 4 but is nonconstructive.

Nash gave 2 different proofs of his existence theorem, one using the Brouwer fixed point theorem (Nash 1950a) and the other using the Kakutani fixed point theorem, a weaker form of the Kakutani-Fan-Glicksberg theorem applicable only to Euclidean spaces (Nash 1950b). Of these, only the second appears to be amenable to generalising as we have done. The Brouwer, Kakutani and Nash theorems are all known to not be provable constructively but they all have equivalent approximation theorems that are provable constructively and are complete for the same complexity class (Tanaka 2011, Daskalakis et al. 2006). The author plans to investigate the computation of abstract Nash equilibria in subsequent papers.

It should be noted that theorem 3, like Nash's original theorem, has a simpler proof using the weaker Kakutani theorem. The reason for proving the stronger theorem 2 is that it can be used to prove a stronger result which also generalises Glicksberg's theorem (Glicksberg 1952), a result which generalises Nash's theorem for finite games to games whose strategy sets are compact topological spaces and whose outcome functions are continuous.

References

- S. Berardi, M. Bezem, and T. Coquand. On the computational content of the axiom of choice. *The journal of symbolic logic*, 63(2):600–622, 1998.
- Ulrich Berger. The Berardi-Bezem-Coquand functional in a domain-theoretic setting. 2002.
- John B. Conway. *A course in functional analysis*. Springer, 1990.
- Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *STOC '06 Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 71–79, 2006.

- Martin Escardó and Paulo Oliva. Selection functions, bar recursion and backward induction. *Mathematical structures in computer science*, 20(2):127–168, 2010a.
- Martin Escardó and Paulo Oliva. What sequential games, the Tychonoff theorem and the double-negation shift have in common. *Proc. of 3rd ACM SIG-PLAN Wksh. on Mathematically Structured Functional Programming*, 2010b.
- Martin Escardó and Paulo Oliva. Sequential games and optimal strategies. *Proc R Soc A*, 2011.
- Martin Escardó and Paulo Oliva. Computing Nash equilibria of unbounded games. *Proceedings of the Turing centenary conference*, 2012.
- Martin Escardó, Paulo Oliva, and Thomas Powell. System T and the product of selection functions. *Computer Science Logic (CSL'11) - 25th International Workshop/20th Annual Conference of the EACSL*, 12:233–247, 2011.
- Ky Fan. Fixed-point and minimax theorems in locally convex topological linear spaces. *Proc Natl Acad Sci U S A*, 38(2):121–126, February 1952.
- Drew Fudenberg and Jean Tirole. *Game theory*. MIT Press, 1991.
- Irving Glicksberg. A further generalisation of the Kakutani fixed point theorem, with application to nash equilibrium points. *Proc. Amer. Math. Soc.*, 3(1): 170–174, February 1952.
- John L. Kelley. *General Topology*. Ishi Press International, 1955.
- John Nash. *Non-cooperative games*. PhD thesis, Princeton University, 1950a.
- John Nash. Equilibrium points in n-person games. *Proc Natl Acad Sci U S A*, 36(1):48–49, January 1950b.
- Thomas Powell. *Ph.d thesis (unpublished)*. PhD thesis, 2012.
- Yasuhito Tanaka. Proof of constructive version of the Fan-Glicksberg fixed point theorem directly by sperner's lemma and approximate nash equilibrium with continuous strategies: a constructive analysis. *IAEng Journal of Applied Mathematics*, 2011.
- John von Neumann and Oskar Morgenstern. *Theory of games and economic behaviour*. Princeton university press, 1944.